

Exploiting Channel Correlation and PU Traffic Memory for Opportunistic Spectrum Scheduling

Shanshan Wang, Sugumar Murugesan and Junshan Zhang

School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ, 85287

Abstract

We consider a cognitive radio network with multiple primary users (PUs) and one secondary user (SU), where a spectrum server is utilized for spectrum sensing and scheduling the SU to transmit over one of the PU channels opportunistically. One practical yet challenging scenario is when *both* the PU occupancy and the channel fading vary over time and exhibit temporal correlations. Little work has been done for exploiting such temporal memory in the channel fading and the PU occupancy simultaneously for opportunistic spectrum scheduling. A main goal of this work is to understand the intricate tradeoffs resulting from the interactions of the two sets of system states - the channel fading and the PU occupancy, by casting the problem as a partially observable Markov decision process. We first show that a simple greedy policy is optimal in some special cases. To build a clear understanding of the tradeoffs, we then introduce a full-observation genie-aided system, where the spectrum server collects channel fading states from all PU channels. The genie-aided system is used to decompose the tradeoffs in the original system into multiple tiers, which are examined progressively. Numerical examples indicate that the optimal scheduler in the original system, with observation on the scheduled channel only, achieves a performance very close to the genie-aided system. Further, as expected, the optimal policy in the original system significantly outperforms randomized scheduling, pointing to the merit of exploiting the temporal correlation structure in both channel fading and PU occupancy.

Index Terms

Cognitive radio networks; spectrum server; temporal correlation; partially observable Markov decision processes; genie-aided system.

I. INTRODUCTION

Over the past decade, cognitive radio (CR) has been identified as one promising solution to ease the “spectrum scarcity” associated with the traditional *static* spectrum allocation [1]–[4]. Going beyond the fixed and licensed spectrum allocation, a secondary user (SU) can opportunistically access the spectrum owned by the primary users (PUs) in a CR network. This paradigm shift from static to dynamic spectrum allocation has been shown to bring significant improvement in the spectrum utilization, and hence the system’s overall performance.

A fundamental principle enabling the cognitive capability is built upon the SU’s dynamic adaptation of its operation parameters (such as power, frequency, etc.), according to the environmental variations over time. One such variation is the channel fading. Often times an *i.i.d.* flat fading model is used in abstracting fading channels, which fails to capture the temporal channel memory observed in realistic scenarios [5]. An alternative model, namely the Gilbert-Elliott (GE) model [6], has been widely used (see, e.g., [7]–[9]) to capture the temporal correlation in the fading process. Specifically, the GE model uses a first-order Markov chain with two states: one representing a “good” channel where the user experiences error-free transmissions, and the other representing a “bad” channel with unsuccessful transmissions.

Another variation to be considered is the PU’s activity on the channels. Note that in a CR network, the SUs have a strictly lower priority in the spectrum usage, and can only access the channels when the PUs are absent [1]–[4]. This unique spectrum usage structure necessitates the inclusion of the channel’s PU occupancy state in determining the channel’s accessibility by an SU.

In many of the existing works (see, e.g., [7], [8], [10]–[12]), only one set of the system states – either the channel fading, or the PU occupancy – has been taken into consideration in developing spectrum access strategies by the SU. In this work, we take a step forward, and explore the utility of *both* the states for opportunistic channel access. Specifically, we consider a CR network consisting of multiple PU channels and one SU opportunistically accessing one of the PU channels at a time. A spectrum server is utilized to periodically schedule the SU to one of the channels for transmission. Worth noting is that the usage of the spectrum server is consistent with the recent FCC ruling on the use of a spectrum database in CR network operations [13]. Further, the spectrum server facilitates spectrum management, and enhances the scalability of the network [14].

Dynamic spectrum access in the presence of the temporal variations can be cast as a sequential control problem. We formulate this sequential control problem as a partially observable Markov decision process [15]. In this context, the spectrum server makes scheduling decisions in terms of allocating a PU channel to the SU, based on the channel’s PU occupancy state and fading state. We model the channel fading by using a two-state first-order Markov chain, i.e., the GE model. On the other hand, since the PU activity may possess a long temporal memory (see, e.g., [16], [17]), we develop an “age” model to capture the temporal correlation structure of the PU occupancy state.

Building on the above model, we examine the intricate tradeoffs resulting from the dynamic interaction of the system states. Our main contributions can be summarized as follows:

- We study opportunistic spectrum scheduling by exploiting the temporal correlation structure in *both* the channel fading and the PU occupancy states. This, to the best of our knowledge, has not been addressed systematically in the literature so far.
- We show that the optimal scheduling involves a *multi-tier* “exploitation vs. exploration” tradeoff. For certain special cases, we establish the optimality of a simple greedy policy, and examine the intricacy of the fundamental tradeoffs.
- To gain a better understanding of the tradeoffs for the general case, we introduce a full-observation genie-aided system, where the spectrum server collects channel fading states from all the PU channels. Using the genie-aided system, we decompose the multiple tiers of the tradeoffs, and examine them progressively.

The rest of the paper is organized as follows. Section II introduces the basic setting and problem formulation in detail. In Section III, we identify the fundamental tradeoffs and illustrate them via special cases. Section IV further examines the tradeoffs by developing a genie-aided system that isolates the impact of channel fading and PU occupancy on the optimal reward. In Section V, numerical results are presented where we evaluate and compare the performance of the optimal policy in the original system with baseline cases. We also study the impact of the memory in the channel fading and PU occupancy on the relative performances of various baseline cases. This is followed by concluding remarks in Section VI.

II. PROBLEM FORMULATION

A. Basic Setting

We consider a CR network with one SU and N PUs¹. Each PU is licensed to one of N independent channels, henceforth identified as PU channels. A PU generates packets according to a stationary process, transmits over its channel if there are backlogged packets, and leaves upon the completion of the transmissions. The PU traffic activity is assumed to be identical and independent across channels.

The SU, on the other hand, is backlogged with packets and opportunistically transmits these packets over the PU channels with the help of a spectrum server. Time is divided into two timescales: mini-slots and the control slots each constituting K mini-slots, as illustrated in Fig. 1. The length of each mini-slot is normalized to fit the transmission of one data packet of the PU or the SU. At the beginning of each control slot, the spectrum server schedules the SU to the “best” PU channel that is expected to yield the highest average throughput for the SU. The SU then transmits packets in the scheduled channel, until it detects the return of a PU². Upon such an event, the SU suspends transmissions until the beginning of the next control slot, when the spectrum server re-schedules the SU to a PU channel based on most recent observations. At the end of each mini-slot when the SU transmitted a packet, it sends accurate feedback on the channel fading state (of the PU channel, as seen by the SU) corresponding to that mini-slot, to the spectrum server. The spectrum server uses this channel fading feedback, the PU traffic observations, along with the memory inherent in these processes to perform informed scheduling decisions at the beginning of the next control slot. We discuss the system model and the scheduling problem formulation in more detail in the following.

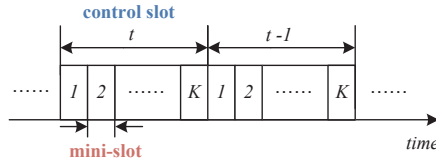


Fig. 1. A sketch of the two timescale model.

B. Problem Formulation

The opportunistic spectrum access at hand can be viewed as a sequential control problem, which we formulate as a partially observable Markov decision process. In the following, we introduce and elaborate the entities involved in the formulation.

Channel occupancy: The usage pattern on each of the PU channels can be modeled as an ON-OFF process at the mini-slot timescale, with ON denoting the *busy* state where the PU transmits data over the channel, and OFF the *idle* state where the PU is absent. Channel occupancy is the idle or busy state of the PU channels. Let $o_{t,k}(n)$ be a binary random variable, denoting

¹Each user is assumed to be a pair of transmitter and receiver.

²This can be accomplished by incorporating collision detection by the SU at the mini-slot timescale. We also assume that PU arrivals coincide with the mini-slot boundaries.

whether PU channel n , for $n \in \{1, \dots, N\}$, is idle ($o_{t,k}(n) = 0$), or not ($o_{t,k}(n) = 1$), in the k th mini-slot of control slot t . The corresponding idle probability is denoted by $\pi_{t,k}^o(n) \triangleq \Pr(o_{t,k}(n) = 0)$.

Idle/Busy age: The PU traffic is temporally correlated, i.e., the current occupancy state on each of the channel depends on the history of the channel occupancy. We introduce the notion of “age,” defined as follows, to characterize the occupancy history:

The age of a PU channel is the number of consecutive mini-slots immediately preceding the current mini-slot, during which the channel is in the same occupancy state as in the current mini-slot. The age is denoted as “idle age” if the channel is in idle state in the current mini-slot and “busy age” otherwise. We use $x_t(n)$ to denote the age of channel n at the beginning of control slot t .

As noted earlier, we assume long memory in the PU occupancy state. Specifically, with the definition of age in place, we adopt a family of functions monotonically decreasing in age, to denote the *conditional probability that a channel will be idle (or busy), given that it has been idle (or busy) for $x \geq 1$ mini-slots:*

$$\begin{aligned} P_I(x) &= \frac{1}{x^u + C_I}, \\ P_B(x) &= \frac{1}{x^u + C_B}, \quad u = 1, 2, \dots, \end{aligned} \quad (1)$$

where C_I and C_B are normalizing constants taking positive values. Our occupancy model essentially imposes the following realistic correlation structures: 1) the occupancy memory weakens with time, i.e., the impact of past occupancy events on the current occupancy state diminishes since the said event happened; 2) the conditional probability that the PU channel is busy or idle now, is purely a function of the length of time the channel has been in the most recent state, and is independent of the channel occupancy history before the time of the latest transition to the most recent state. In sight of this, the quantities P_I and P_B defined in (1) are sufficient for capturing the temporal correlation in the channels’ PU occupancy state.

Channel fading model: At the end of each mini-slot after transmitting a packet, the SU measures the channel fading between its transmitter and receiver on the scheduled channel, and feeds back this information to the spectrum server. Inspired by recent works [7], [8], [10], we capture the memory in the fading (of the PU channel) between the SU’s transmitter and receiver using a two-state, first-order Markov chain, with state variations occurring at the mini-slot timescale. The Markov chain model is *i.i.d.* across the PU channels. Each state of the Markov chain corresponds to the degree of decodability of the data sent through the channel, where state 1 denotes full decodability and state 0 denotes zero decodability. Note that the states can also be interpreted as a quantized representation of the underlying channel fading, in the sense that state 1 corresponds to “good” channel fading, while state 0 corresponds to “bad” fading. The probability transition matrix of this Markov chain is given as:

$$\mathbf{P} := \begin{bmatrix} 1-r & r \\ 1-p & p \end{bmatrix}, \quad (2)$$

where p is the conditional probability that the channel fading is good, given that it was good in the previous mini-slot; and r is the conditional probability that the channel fading is good, given that it was bad in the previous mini-slot. Throughout the paper, we will focus on the case when the fading channels are positively correlated, i.e., $p > r$.

Belief of channel fading state: Denote by $\pi_{t,k}^s(n)$ the belief of channel fading state in the k th mini-slot of control slot t on channel n . As is standard [7], [15], the fading state belief is a sufficient statistic that characterizes the current channel fading state as perceived by the SU. Further, let $f_{t,k}(a_t)$ be a binary random variable denoting the fading state feedback obtained at the end of the k th mini-slot in control slot t on the scheduled channel a_t . Also, define $T^L(\cdot)$, for $L \in \{0, 1, \dots\}$, as the L th step belief evolution operator, taking the form: for $\gamma \in (0, 1)$,

$$T^L(\gamma) = T(T^{L-1}(\gamma)), \quad (3)$$

with $T^0(\gamma) = \gamma$ and $T(\gamma) = \gamma p + (1 - \gamma)r$. Now, the update of the fading state belief is governed by the underlying Markov chain model, and any new information obtained on the channel fading, i.e.:

$$\pi_{t,k+1}^s(n) = \begin{cases} p, & a_t = n, f_{t,k}(a_t) = 1, \\ r, & a_t = n, f_{t,k}(a_t) = 0, \\ T(\pi_{t,k}^s(n)), & a_t \neq n. \end{cases} \quad (4)$$

Action space: This refers to the set of channels that the scheduling decision is made from. The spectrum server selects channels only from those that are currently idle³, and the action space \mathcal{A}_t in control slot t can thus be written as:

$$\mathcal{A}_t = \{n : o_{t,1}(n) = 0\}. \quad (5)$$

³This is a policy level constraint to protect the PU’s priority in spectrum access.

State: At the beginning of each control slot, the spectrum server makes the scheduling decision based on three factors: For each of the PU channels, 1) the idle/busy state at the moment; 2) the length of time the channel has been in the current occupancy state (i.e., age); and 3) the fading state belief value. That is, the state of each PU channel n , is represented by a three dimensional vector: $S_t(n) = [o_{t,1}(n), x_t(n), \pi_{t,1}^s(n)]$. Accordingly, the state of the system at the beginning of current control slot t is described by a $N \times 3$ matrix \mathbf{S}_t :

$$\mathbf{S}_t := [S_t(1); \dots; S_t(N)] = \begin{bmatrix} o_{t,1}(1) & x_t(1) & \pi_{t,1}^s(1) \\ o_{t,1}(2) & x_t(2) & \pi_{t,1}^s(2) \\ \vdots & \vdots & \vdots \\ o_{t,1}(N) & x_t(N) & \pi_{t,1}^s(N) \end{bmatrix}. \quad (6)$$

Horizon: The horizon is the number of consecutive control slots over which scheduling is performed. We index the control slots in a decreasing order with control slot 1 being the end of the horizon⁴. Throughout the paper, we denote the length of the horizon by m , i.e., the scheduling process begins at control slot m .

Stationary scheduling policy: A stationary scheduling policy \mathcal{P} establishes a stationary mapping from the current state \mathbf{S}_t to an action a_t in each control slot t .

Expected immediate reward: The expected immediate reward is the reward accrued by the SU within the current control slot. Specifically, the SU collects one unit of reward in each mini-slot, if the channel is idle and has good channel fading (i.e., conditions that indicate successful transmission by SU). Since the scheduled channel must be idle in the first mini-slot of the current control slot, the expected immediate reward can be calculated as:

$$R_t(\mathbf{S}_t, a_t) = \sum_{k=2}^K \pi_{t,k}^o(a_t) \pi_{t,k}^s(a_t) + \pi_{t,1}^s(a_t). \quad (7)$$

Total discounted reward: Given a scheduling policy \mathcal{P} , the total discounted reward, accumulated from the current control slot t , until the horizon, can be written as⁵

$$V_t(\mathbf{S}_t; \mathcal{P}) = R_t(\mathbf{S}_t, a_t) + \beta E_{\pi_{t-1}^s} E_{\mathbf{O}_{t-1}} V_{t-1}(\mathbf{S}_{t-1}; \mathcal{P}), \quad (8)$$

where $\beta \in (0, 1)$ is the discount factor, facilitating relative weighing between the immediate and future rewards, and the expectation is taken with respect to fading state belief: $\pi_{t-1}^s = \{\pi_{t-1,1}^s(1), \dots, \pi_{t-1,1}^s(N)\}$, and PU occupancy: $\mathbf{O}_{t-1} = \{o_{t-1,1}(1), \dots, o_{t-1,1}(N)\}$.

Objective function: The objective of the scheduling problem is to maximize the SU's throughput, i.e., SU's total discounted reward. A scheduling policy \mathcal{P}^* is optimal if and only if the following optimality equation is satisfied:

$$V(\mathbf{S}_t; \mathcal{P}^*) = \max_{a_t \in \mathcal{A}_t} \left\{ R_t(\mathbf{S}_t, a_t) + \beta E_{\pi_{t-1}^s} E_{\mathbf{O}_{t-1}} V(\mathbf{S}_{t-1}; \mathcal{P}^*) \right\}. \quad (9)$$

The function $V^*(\mathbf{S}_t) := V(\mathbf{S}_t; \mathcal{P}^*)$ is the objective function of the scheduling problem.

III. FUNDAMENTAL TRADEOFFS

The decision on opportunistic spectrum scheduling is made based on two sets of system states: the PU occupancy on the channel and the channel fading perceived by the SU. On one hand, PUs may return in the middle of a control slot and hinder further transmissions of the SU, leading to a decreased reward for the SU. The temporal memory resident in the PU occupancy suggests that the past history of channel's occupancy, measured by the age, influences the occupancy state of the channel in the future. On the other hand, the PU channels may suffer from "bad" channel fading in the middle of a control slot, even if a PU does not return to hinder SU's transmissions.. Similar to the PU occupancy, the historic observation on the fading process would help determine the expected channel fading in the future. Note that by way of the channel feedback arrangement, an observation of a PU channel fading is made only when that channel is scheduled to the SU. Thus scheduling is inherently tied to channel fading learning. Roughly speaking, to maximize the SU's reward, the spectrum server must schedule a channel such that the combination of the perceived channel occupancy and channel fading strikes a "perfect" balance between the immediate gains and channel learning for future gains. We discuss this intricate tradeoff in the following.

A. Classic "Exploitation vs. Exploration" Tradeoff

In the existing literature (e.g., [7], [8], [10]–[12]), focus has been cast on considering only one of the factors: either channel fading or channel occupancy, along with the associated temporal correlation. The optimal decision is a mapping that best

⁴For the mini-slots, we use the conventional increasing time indexing.

⁵In the subsequent sections, we may drop \mathcal{P} and the tiers of expectation to simplify the notation.

balances the tradeoff of “exploitation” and “exploration” on the single factor being considered. The exploitation side lets the scheduler choose the channel with the best perceived channel fading (or occupancy state) at the moment, corresponding to immediate gains; while the exploration side tends to favor the channel with the least learnt information so far, probing which can contribute to the overall understanding of the channel fading (or occupancy state) in the network, and thus better opportunistic scheduling decisions in the future.

B. “Exploitation vs. Exploration” Tradeoff in Dynamics of Both Channel Fading and PU Occupancy States

In contrast to the existing works, we examine the tradeoffs when the temporal correlation in *both* the channel fading and PU occupancy are considered. While the classic tradeoff described above apparently exists, additional tradeoffs arise in our context due to the interactions between the two sets of system states. In particular, the long temporal memory in the PU occupancy state adds a new layer to the tradeoffs inherent in the problem. For instance, note that the SU can only transmit on *idle* channels. To carry out exploration, the channel being favored in the traditional sense, i.e., the one on which the least information is available, may no longer be the preferred choice if this channel is perceived to be unavailable (i.e., busy) for a prolonged duration in the near future. In other words, it may not be worth learning the channel as the SU cannot utilize the learned knowledge in the near future.

Fig. 2 is a pictorial illustration of the impact from the occupancy state on the SU’s expected reward. The history of occupancy, represented by the idle ages $x_t(1)$ and $x_t(2)$, affects both the immediate and future rewards of the SU. Specifically, as the idle age increases, the temporal memory in the PU’s occupancy pushes the channel to transit to busy sooner (i.e., time point b comes earlier than a in the figure). Therefore, the average availability on the PU channel in the current control slot decreases, which leads to a smaller immediate reward for the SU.

Further, note that the latest mini-slot for which the spectrum server receives channel fading feedback is also the last mini-slot before the PU returns, i.e., time points a and b respectively for the two cases in Fig. 2. The duration d_1 (likewise, d_2) in Fig. 2 is an indication of how “fresh” the channel fading information is for the scheduling decision at the beginning of the next control slot, i.e., $t - 1$. With $d_1 < d_2$, channel feedback is more fresh in the former case, with a lower idle age $x_t(1)$. Thus, age, through its effect on the *freshness* of feedback, and the availability of the PU channels in the future slots, adds another layer to the tradeoffs, thereby influencing the optimal scheduling decision.

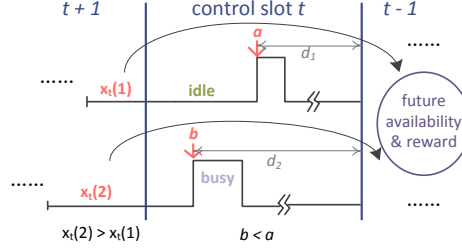


Fig. 2. An illustration of the impact of age on the SU’s reward.

To better perceive the intricate tradeoffs in the system, we proceed, in what follows, with a number of break-down results that aim at illustrating each tier of the tradeoffs progressively.

C. Tradeoffs Inherent in Immediate Reward

Consider the PU channel scheduled in the current control slot t . Let k_0 denote the latest mini-slot before the PU of the scheduled channel returns in the current control slot. Clearly, k_0 is a random variable, taking values in $k_0 \in \{1, \dots, K\}$. Let $p_z \triangleq \Pr(k_0 = z)$. With x denoting the idle age of the scheduled channel, k_0 is distributed as follows: For $K = 2$:

$$p_1 = 1 - P_I(x + 1), \quad p_2 = P_I(x + 1),$$

and for $K \geq 3$:

$$\begin{aligned} p_1 &= 1 - P_I(x + 1), \\ p_z &= \prod_{j=1}^{z-1} P_I(x + j)(1 - P_I(x + z)), \quad z = 2, \dots, K - 1, \\ p_K &= \prod_{j=1}^{K-1} P_I(x + j). \end{aligned} \tag{10}$$

In the following lemma, we establish the structure of the distribution of k_0 .

Lemma 1: The distribution of k_0 is monotonically decreasing in the idle age, x_t , for $z = 2, \dots, K$, and monotonically increasing in x_t , for $z = 1$.

Proof: For $K = 2$, it is straightforward to establish the conclusion. We next focus on $K \geq 3$. First, it is easy to show that for $z = 1$, p_1 is monotonically increasing in x since $P_I(x)$ is monotonically decreasing. Further, note that if two positively valued functions, $f_1(x) > 0$ and $f_2(x) > 0$, are both monotonically decreasing in x , i.e., for any positive integer $\Delta \geq 1$,

$$f_1(x) > f_1(x + \Delta), \quad f_2(x) > f_2(x + \Delta),$$

then the product of the two, $f_1(x)f_2(x)$, is also a decreasing function since

$$f_1(x)f_2(x) > f_1(x + \Delta)f_2(x + \Delta).$$

Therefore, p_K is monotonically decreasing in x .

Next, we show that $p_z, z = 2, \dots, K - 1$ are monotonically decreasing in x . Based on the above argument, it is sufficient to show that for any $z = 2, \dots, K - 1$, the following function,

$$g(x) \triangleq P_I(x + z - 1)(1 - P_I(x + z)),$$

is monotonically decreasing. We have the following simplifications: For $\Delta \geq 1$, a positive integer,

$$\begin{aligned} \frac{g(x)}{g(x + \Delta)} &= \frac{P_I(x + z - 1)(1 - P_I(x + z))}{P_I(x + \Delta + z - 1)(1 - P_I(x + \Delta + z))} \\ &= \frac{(x + \Delta + z - 1)^u + C_I}{(x + z - 1)^u + C_I} \cdot \frac{(x + \Delta + z)^u + C_I - 1}{(x + \Delta + z)^u + C_I - 1} \cdot \frac{(x + z)^u + C_I - 1}{(x + z - 1)^u + C_I} \\ &\triangleq B_1 B_2 B_3. \end{aligned} \tag{11}$$

Since $x \geq 0, C_I > 0, u \geq 1$ and $\Delta \geq 1$, we immediately have $B_1 \geq 1, B_2 > 1$. Further, using the binomial theorem, we obtain

$$(x + z)^u = \sum_{i=0}^u \binom{u}{i} (x + z - 1)^i \geq (x + z - 1)^u + 1, \tag{12}$$

and hence $B_3 \geq 1$. As a result, $\frac{g(x)}{g(x + \Delta)} > 1$ and p_2, \dots, p_{K-1} are monotonically decreasing in x . This concludes the proof. ■

We next present a result that demonstrates the tradeoff inherent in the immediate reward with respect to age.

Proposition 1: The immediate reward on the scheduled channel is monotonically decreasing in the idle age.

Proof: As the system model implies, we can rewrite the immediate reward as the following weighted sum:

$$R_t(\mathbf{S}_t, a_t) = \sum_{z=1}^K \sum_{k=1}^z \pi_{t,k}^s(a_t) p_z, \tag{13}$$

where $p_z = \Pr(k_0 = z)$ is given by (10).

When $K = 2$, with idle age on a_t being x_t , we have $R_t(\mathbf{S}_t, a_t) = \pi_{t,1}^s(a_t) + P_I(x_t + 1)\pi_{t,2}^s(a_t)$. Apparently, it increases as x_t decreases.

For $K \geq 3$, denote by $\theta_z \triangleq \sum_{k=1}^z \pi_{t,k}^s(a_t), z = 1, \dots, K$. It is clear that

$$\theta_1 < \theta_2 < \dots < \theta_K, \tag{14}$$

and $\{\theta_z\}$'s are constants in the idle age x_t . To emphasize the role of the argument x_t , we rewrite $R_t(\mathbf{S}_t, a_t)$ as $R_t(x_t)$, and p_z as $p_z(x_t)$.

Utilizing the result of Lemma 1, and noting that for any positive integer $\Delta \geq 1$, $|p_1(x_t) - p_1(x_t + \Delta)| = \sum_{z=2}^K (p_z(x_t) - p_z(x_t + \Delta))$, we obtain:

$$R_t(x_t) - R_t(x_t + \Delta) = \sum_{z=1}^K \theta_z (p_z(x_t) - p_z(x_t + \Delta)) > \theta_1 (p_1(x_t) - p_1(x_t + \Delta)) + \theta_2 \sum_{z=2}^K (p_z(x_t) - p_z(x_t + \Delta)) > 0,$$

i.e., $R_t(x_t)$ is monotonically decreasing in the idle age x_t , and this establishes the proposition. ■

The above result can be readily extended to the following corollary.

Corollary 1: When all the PU channels have equal fading state beliefs, the immediate reward is maximized by scheduling the SU to the channel with the lowest idle age.

Next, recall that the channel fading is modeled by a positively-correlated Markov chain. Hence, if $\pi_{t,1}^s(a_t) > \pi_{t,1}^{s'}(a_t)$, then the inequality $\pi_{t,k}^s(a_t) > \pi_{t,k}^{s'}(a_t)$ holds, for all $k = 2, \dots, K$. We present the following proposition without further proof.

Proposition 2: The immediate reward on the scheduled channel is monotonically increasing in its fading state belief at the moment. Further, given equal idle ages across all PU channels, the immediate reward is maximized by scheduling the SU to the channel with the largest fading state belief value at the moment.

D. Tradeoffs Inherent in Total Reward

In this subsection, we illustrate the tradeoffs inherent in the total reward by examining a special case with two channels $N = 2$ and number of mini-slots $K = 1$. In particular, we show that under these conditions, a simple greedy scheduling policy is optimal. The greedy policy is formally defined as follows: In any control slot, the greedy decision maximizes the immediate reward, ignoring the future rewards, i.e.,

$$\hat{a}_t = \max_{a_t \in \mathcal{A}_t} \{R_t(\mathbf{S}_t, a_t)\}. \quad (15)$$

We now formally record the result on greedy policy optimality in the following proposition.

Proposition 3: The greedy policy is optimal when $K = 1$ and $N = 2$.

Proof: To prove the proposition, we begin with the following induction hypothesis:

Induction Hypothesis: With the length of the horizon denoted by $m, m \geq 2$, assume that greedy policy is optimal in all the control slots $t \in \{m-1, \dots, 1\}$.

The proof proceeds in two steps: First, we fix a sequence of scheduling decisions $\vec{l} := \{a_m, \dots, a_{t+1}\}$, and show that the expected immediate reward in control slot t , under the greedy policy, is independent of the scheduling decisions \vec{l} . Then, we provide induction based arguments to validate the induction hypothesis and hence establish that the greedy policy is optimal in all the control slots.

Let $U_t^{(\vec{l})}$ denote the expected immediate reward in slot $t \in \{m-1, \dots, 1\}$, given the scheduling decisions \vec{l} . $U_t^{(\vec{l})}$ can be calculated as:

$$U_t^{(\vec{l})} = \sum_{\{o_t(1), o_t(2)\}} U_t^{(\vec{l})}(o_t(1), o_t(2)) \Pr(o_t(1), o_t(2)), \quad (16)$$

where $o_t(1)$ (likewise, $o_t(2)$) is the binary indicator of whether channel 1 (or 2) is idle ($o_t(1) = 0$) or not ($o_t(1) = 1$) in the t th control slot, and $\Pr(o_t(1), o_t(2))$ denotes the joint probability of both channels' availability status (idle or busy). There exist four realizations of the vector $(o_t(1), o_t(2))$, namely $\{(0, 0), (0, 1), (1, 0), (1, 1)\}$. In what follows, we show that the value of $U_t^{(\vec{l})}(o_t(1), o_t(2))$ calculated under each of the realizations is independent of the scheduling decisions \vec{l} .

Case 1: When $(o_t(1), o_t(2)) = (0, 0)$. In this case, both channels are idle in control slot t . Let $\pi_{t+1}^s(n)$ be the fading state belief on channel n in control slot $t+1$. The expected immediate reward in control slot t , under the greedy policy, can be calculated as

$$U_t^{(\vec{l})}(0, 0) = \pi_{t+1}^s(a_{t+1})p + (1 - \pi_{t+1}^s(a_{t+1}))T(\pi_{t+1}^s(\tilde{a}_{t+1})), \quad (17)$$

where $\tilde{a}_{t+1} = \{1, 2\} \setminus a_{t+1}$. For notational convenience, we write $\alpha \triangleq \pi_{t+1}^s(a_{t+1})$ and $\tilde{\alpha} \triangleq \pi_{t+1}^s(\tilde{a}_{t+1})$. The reward $U_t^{(\vec{l})}(0, 0)$ can now be further expressed as

$$\begin{aligned} U_t^{(\vec{l})}(0, 0) &= p\alpha + (1 - \alpha)T(\tilde{\alpha}) \\ &= p\alpha + (1 - \alpha)(\alpha p + (1 - \alpha)r) \\ &= pP_1 + rP_2, \end{aligned} \quad (18)$$

where

$$P_1 \triangleq \alpha + \tilde{\alpha} - \alpha\tilde{\alpha}, \quad P_2 \triangleq (1 - \alpha)(1 - \tilde{\alpha}). \quad (19)$$

That is, P_1 is the probability that at least one of the channels experiences good channel fading in the previous control slot $t+1$, while P_2 is the probability that both channels see bad channel fading. It is noted that these probabilities are controlled by the underlying Markov dynamics only, and thus P_1 and P_2 are independent of the scheduling decisions \vec{l} . Therefore, $U_t^{(\vec{l})}(0, 0)$ is independent of \vec{l} .

Case 2: When $(o_t(1), o_t(2)) = (0, 1)$. In this case, only channel 1 is idle and can be scheduled. The reward $U_t^{(\vec{l})}(0, 1)$ is obtained as

$$U_t^{(\vec{l})}(0, 1) = \begin{cases} p\pi_{t+1}^s(1) + r(1 - \pi_{t+1}^s(1)), & \text{if } a_{t+1} = 1, \\ T(\pi_{t+1}^s(1)), & \text{if } a_{t+1} = 2. \end{cases} \quad (20)$$

It follows from (3) that

$$U_t^{(\vec{l})}(0, 1)|_{a_{t+1}=1} = U_t^{(\vec{l})}(0, 1)|_{a_{t+1}=2}, \quad (21)$$

i.e., $U_t^{(\vec{l})}(0, 1)$ is independent of \vec{l} .

Case 3: When $(o_t(1), o_t(2)) = (1, 0)$. Similar to *Case 2*, only channel 2 can be scheduled in this case, and we have:

$$U_t^{(\vec{l})}(1, 0) = \begin{cases} T(\pi_{t+1}^s(2)), & \text{if } a_{t+1} = 1, \\ p\pi_{t+1}^s(2) + r(1 - \pi_{t+1}^s(2)), & \text{if } a_{t+1} = 2. \end{cases} \quad (22)$$

Again, this indicates that

$$U_t^{(\vec{l})}(1, 0)|_{a_{t+1}=1} = U_t^{(\vec{l})}(1, 0)|_{a_{t+1}=2}, \quad (23)$$

i.e., $U_t^{(\vec{l})}(1, 0)$ is independent of \vec{l} .

Case 4: When $(o_t(1), o_t(2)) = (1, 1)$. In this case, both channels are busy, and it follows immediately that

$$U_t^{(\vec{l})}(1, 1) = 0. \quad (24)$$

Clearly, $U_t^{(\vec{l})}(1, 1)$ is independent of \vec{l} as well.

Next, note that $\Pr(o_t(1), o_t(2))$ is a function of the ages $(x_t(1), x_t(2))$ only, which evolve independently from the scheduling decisions \vec{l} . Thereby, we conclude that $\Pr(o_t(1), o_t(2))$ is independent of the scheduling decisions \vec{l} , and so is the expected immediate reward in control slot t , i.e., $U_t^{\vec{l}} = U_t$. Now, by extension, we have that the total reward collected from control slot t till the horizon is independent of \vec{l} , i.e.,

$$\sum_{t'=t}^1 U_{t'}^{\vec{l}} = \sum_{t'=t}^1 U_{t'}.$$

Thus, the greedy policy is optimal in control slot $t+1$ as well. Since $t \in \{m-1, \dots\}$ is arbitrary, the greedy policy is optimal in every control slot $\{m, \dots, 1\}$ under the induction hypothesis. Finally, as the greedy policy is trivially optimal at the horizon, i.e., $t = 1$, using backward induction, we validate the induction hypothesis, and arrive at the conclusion that greedy is optimal in all control slots $t \in \{m, \dots, 1\}$. This establishes the proposition. ■

Remarks: Note that the tradeoffs inherent in the special case considered above, i.e., $K = 1, N = 2$, is more intricate than those observed in related recent works (e.g., [7], [8]), where a control slot coincides with a mini-slot and only one of the states: channel fading or PU occupancy, is considered. This is because, despite $K = 1$, the question of “whether to learn a channel that may not be available for scheduling in the near future due to channel occupancy state” still exists. Thus the tradeoffs discussed in the preceding subsections are retained in this special case, essentially adding value to our result on greedy optimality. In the subsequent section, we proceed to further understand the tradeoffs in the original system by introducing a conceptual “genie-aided system.”

IV. MULTI-TIER TRADEOFFS: A CLOSER LOOK VIA A GENIE-AIDED SYSTEM

In the previous section, we partially showed the interaction between various state elements by examining the immediate reward and certain special cases. In order to obtain a more complete understanding of the inherent dynamics, we next introduce a full-observation genie-aided system that helps decompose and characterize the various tiers of the multi-dimensional tradeoffs.

A. A Genie-Aided System

The genie-aided system is a variant of the original system with the following modification: The spectrum server receives channel fading feedbacks from *all* the channels and not only the scheduled channel. These feedbacks are collected at the same times as those of the feedback from the scheduled channel. Thus when the PU returns on the scheduled channel, the feedback from *all* the channels stop at once. Note that this is a conceptual system, without practical significance, which as we will see, is helpful in better understanding the complicated tradeoffs inherent in the original system.

B. Tradeoffs Associated with Channel Fading

Proposition 4: When the idle ages are the same across all PU channels, it is optimal to schedule the SU to the channel with the highest fading state belief at the moment, i.e.,

$$a_t^* = \arg \max_n \{\pi_{t,1}^s(n)\}. \quad (25)$$

Proof: First, from Proposition 2, the immediate reward is maximized when scheduling the channel with the highest fading state belief at the moment. Now, we focus on showing that the future reward is independent of the action in the current control slot and therefore establishing the proposition. Specifically, at the current control slot t , $t \geq 2$, consider an arbitrary control slot in the future, $t_0 \in \{t-1, \dots, 1\}$. In the following, we show that the expected immediate reward in this control slot, denoted by U_{t_0} , is independent of the current action a_t , and thus the future reward is independent of a_t .

Let k'_0 denote the latest idle mini-slot in control slot $t_0 + 1$ before the PU returns. The reward U_{t_0} is then calculated as:

$$U_{t_0} = \sum_{k'_0=1}^K U_{t_0}(k'_0) \Pr(k'_0), \quad (26)$$

where $\Pr(k'_0)$ is the distribution of k'_0 , identical to that of k_0 as given in (10), and $U_{t_0}(k'_0)$ is the expected reward in control slot t_0 for a given k'_0 .

In the genie-aided system, the spectrum server obtains the feedbacks of the fading states from *all* N channels simultaneously, i.e., at the end of mini-slot k'_0 . Place the channels on which good channel fading is observed at k'_0 in the set \mathcal{C}_1 , and the rest in another set \mathcal{C}_0 . The characterization of the reward $U_{t_0}(k'_0)$ can be further divided into the following cases.

- *Case 1:* $\mathcal{C}_0 = \emptyset$. This corresponds to the case where $f_{t_0+1,k'_0}(n) = 1, \forall n = 1, \dots, N$. Let $W_{t_0}^{(\text{case1})}(n)$ be the expected reward in control slot t_0 on channel n in this case. We have

$$W_{t_0}^{(\text{case1})}(n) = \pi_{t_0,1}^o(n) \sum_{z=1}^K \sum_{k=1}^z \pi_{t_0,k}^s(n) p_z(x_{t_0}(n)), \quad (27)$$

where $\pi_{t_0,1}^s(n) = T^{K-k'_0}(p), \forall n \in \mathcal{C}_1$. It follows that $\pi_{t_0,k}^s(1) = \dots = \pi_{t_0,k}^s(N), \forall k = 1, \dots, K$. Further, since $x_t(1) = \dots = x_t(N)$, we have $p_z(x_{t_0}(1)) = \dots = p_z(x_{t_0}(N))$, and $\pi_{t_0,1}^o(1) = \dots = \pi_{t_0,1}^o(N)$, and therefore, scheduling *any* of the idle channels in \mathcal{C}_1 achieves the same expected reward in control slot t_0 , i.e.,

$$U_{t_0}(k'_0)|_{\text{case1}} = W_{t_0}^{(\text{case1})}(n), \forall n \in \mathcal{C}_1, o_{t_0,1}(n) = 0. \quad (28)$$

Now, note that given equal idle ages on all N channels, the distribution of k'_0 is identical across channels. Therefore, for all $n = 1, \dots, N$, the values of $\pi_{t_0,k}^s(n), k = 1, \dots, K, \pi_{t_0,1}^o(n)$ and $p_z(x_{t_0}(n))$ all stay unchanged if a different channel $a_t \neq a'_t$ is scheduled in the current control slot t . This implies that $U_{t_0}(k'_0)|_{\text{case1}}$ is unchanged, and is thus independent of a_t .

- *Case 2:* $\mathcal{C}_1 = \emptyset$. In this case, $f_{t_0+1,k'_0}(n) = 0, \forall n = 1, \dots, N$. The expected reward collected in control slot t_0 on channel n , denoted by $W_{t_0}^{(\text{case2})}(n)$, can be expressed the same as (27), where the channel strength belief $\pi_{t_0,1}^s(n) = T^{K-k'_0}(r)$, for all $n \in \mathcal{C}_0$. Then, using the similar reasoning as in *Case 1*, we obtain that scheduling *any* of the idle channels in \mathcal{C}_0 achieves the same expected reward in control slot t_0 , i.e.,

$$U_{t_0}(k'_0)|_{\text{case2}} = W_{t_0}^{(\text{case2})}(n), \forall n \in \mathcal{C}_0, o_{t_0,1}(n) = 0. \quad (29)$$

Further, the expected reward achieved in this case, $U_{t_0}(k'_0)|_{\text{case2}}$, does not change when the action in current control slot t varies, i.e., $U_{t_0}(k'_0)|_{\text{case2}}$ is independent of a_t .

- *Case 3:* $\mathcal{C}_0 \neq \emptyset, \mathcal{C}_1 \neq \emptyset$. In this case, we first show that the maximum expected reward in control slot t_0 is achieved by scheduling any of the idle channels in the set \mathcal{C}_1 , which are perceived to have better channel fading state in the subsequent control slot t_0 than the channels in set \mathcal{C}_0 . Specifically, picking any one of the channels from each of the set, $n_1 \in \mathcal{C}_1$ and $n_0 \in \mathcal{C}_0$, we have

$$\begin{aligned} W_{t_0}^{(\text{case3})}(n_1) &= \sum_{z=1}^K \sum_{k=1}^z \pi_{t_0,k}^s(n_1) p_z(x_{t_0}(n_1)), \\ W_{t_0}^{(\text{case3})}(n_0) &= \sum_{z=1}^K \sum_{k=1}^z \pi_{t_0,k}^s(n_0) p_z(x_{t_0}(n_0)), \end{aligned} \quad (30)$$

where

$$\begin{aligned} \pi_{t_0,k}^s(n_1) &= T^{K-k'_0+k-1}(p), \\ \pi_{t_0,k}^s(n_0) &= T^{K-k'_0+k-1}(r). \end{aligned} \quad (31)$$

Based on (3) and the property of the positively-correlated Markov chain, the following inequality holds: For all $k = 1, \dots, K$,

$$\pi_{t_0,k}^s(n_1) = T^{K-k'_0+k-1}(p) \geq T^{K-k'_0+k-1}(r) = \pi_{t_0,k}^s(n_0),$$

with the equality achieved when $K \rightarrow \infty$. Further, applying similar argument as before, we obtain

$$W_{t_0}^{(\text{case3})}(n_1) > W_{t_0}^{(\text{case3})}(n_0).$$

Now, since $n_1 \in \mathcal{C}_1$ (likewise, $n_0 \in \mathcal{C}_0$) is arbitrary, from the conclusion drawn in *Case 1*, the maximal expected reward

in control slot t under this case is achieved by scheduling the SU to *any* of the idle channels in the set \mathcal{C}_1 , i.e.,

$$U_{t_0}(k'_0)|_{\text{case3}} = W_{t_0}^{(\text{case3})}(n_1), \forall n_1 \in \mathcal{C}_1, o_{t_0,1}(n_1) = 0. \quad (32)$$

Next, based on the similar reasoning as in the previous cases, we know that $U_{t_0}(k'_0)|_{\text{case3}}$ is independent of a_t .

Finally, note that $U_{t_0}(k'_0)$ can be written as:

$$U_{t_0}(k'_0) = \sum_{i=1}^3 U_{t_0}(k'_0)|_{\text{case } i} \Pr(\text{Case } i), \quad (33)$$

where

$$\begin{aligned} \Pr(\text{Case1}) &= \prod_{n \in \mathcal{C}_1} \pi_{t_0+1, k'_0}^s(n), \\ \Pr(\text{Case2}) &= \prod_{n \in \mathcal{C}_0} (1 - \pi_{t_0+1, k'_0}^s(n)), \\ \Pr(\text{Case3}) &= \prod_{n \in \mathcal{C}_1} \pi_{t_0+1, k'_0}^s(n) \prod_{n' \in \mathcal{C}_0} (1 - \pi_{t_0+1, k'_0}^s(n')), \end{aligned}$$

are quantities dependent on k'_0 only. Based on the fact that the idle age at the moment are identical across the channels, the probabilities $\Pr(\text{Case } i), i = 1, 2, 3$, and the distribution $\Pr(k'_0)$, are the same across the channels as well. Thus U_{t_0} is independent of a_t .

Since $t_0 \in \{t-1, \dots, 1\}$ is arbitrary, by extension, we have that the total reward collected from control slot t_0 till the horizon, i.e., $\sum_{t'=t_0}^1 U_{t'}$, which is the future reward of current control slot t , is independent from a_t . This concludes the proof and establishes the proposition. ■

Remarks: Proposition 4 illustrates the effect of fading state belief on the optimal decisions in the genie-aided system. With ages equalized across the PU channels and the classic “exploitation vs. exploration” tradeoff neutralized (by definition of the genie-aided system), we saw that, higher fading belief favors the immediate reward and that the future reward is, in fact, independent of the current decision.

In the following, we study the effect of PU occupancy and age on the optimal decisions in the genie-aided system and, in turn, its impact on the original system.

C. Tradeoffs Associated with PU Occupancy

The following proposition identifies the effect of channel occupancy state on the optimal scheduling decisions.

Proposition 5: When the fading state beliefs are the same across all PU channels, it is optimal to schedule the SU to the channel with the lowest idle age at the moment, i.e.,

$$a_t^* = \arg \min_n \{x_t(n)\}. \quad (34)$$

Proof: We prove the proposition by showing that scheduling the channel with the lowest idle age favors: 1) the immediate reward $R_t(\mathbf{S}_t, a_t)$; and 2) the optimal future reward $V_{t-1}^*(\mathbf{S}_{t-1})$. The first part can be readily shown by appealing to Proposition 1 and Corollary 1. To show the second part, we adopt the following induction hypothesis:

Induction Hypothesis: The optimal future reward is convex in the fading state belief.

When channel a_t is scheduled in the current control slot t , the expected future reward can be evaluated as:

$$V_{t-1}^*(\mathbf{S}_{t-1})|_{a_t} = \pi_{t, k_0}^s(a_t) V_{t-1}^*(p) + (1 - \pi_{t, k_0}^s(a_t)) V_{t-1}^*(r),$$

where $V_{t-1}^*(p)$ and $V_{t-1}^*(r)$ represent the future reward calculated when the channel fading state observed in the k_0 th mini-slot of control slot t is good or bad, respectively. More specifically,

$$\begin{aligned} V_{t-1}^*(p) &\triangleq V_{t-1}^*(\pi_{t-1,1}^s(a_t) = \mathbb{T}^{K-k_0}(p)), \\ V_{t-1}^*(r) &\triangleq V_{t-1}^*(\pi_{t-1,1}^s(a_t) = \mathbb{T}^{K-k_0}(r)). \end{aligned}$$

Based on (3), we have, for $\gamma \in (0, 1)$,

$$\mathbb{T}^L(\gamma) = (p-r)^L \gamma + r \frac{1 - (p-r)^L}{1 - (p-r)}, L = 0, 1, \dots,$$

and $\lim_{L \rightarrow \infty} \mathbb{T}^L(\gamma) = \frac{r}{1-p+r} \triangleq \pi_{s|s}$, where $\pi_{s|s}$ denotes the steady-state probability of perceiving good channel fading on any of the PU channels. This indicates that, a smaller k_0 , associated with a higher idle age at the moment (recall discussions

in Section III), results in a larger $K - k_0$ and thus a value of $\pi_{t-1,1}^s(a_t)$ closer to $\pi_{s|s}$, in which case,

$$V_{t-1}^*(\mathbf{S}_{t-1})|_{a_t} \rightarrow \pi_{t,k_0}^s(a_t)V_{t-1}^*(\pi_{s|s}) + (1 - \pi_{t,k_0}^s(a_t))V_{t-1}^*(\pi_{s|s}) \triangleq V_{t-1}^*(a_t, E_{\pi_{s|s}}).$$

On the contrary, as k_0 becomes larger because of a lower idle age, the value of $\pi_{t-1,1}^s(a_t)$ deviates further away from $\pi_{s|s}$, but is closer to p (or r). Also, $\pi_{t,k_0}^s(a_t)$ gets closer to $\pi_{s|s}$. Therefore,

$$V_{t-1}^*(\mathbf{S}_{t-1})|_{a_t} \rightarrow \pi_{s|s}V_{t-1}^*(p) + (1 - \pi_{s|s})V_{t-1}^*(r) \triangleq E_{\pi_{s|s}}V_t^*(a_t).$$

Now, appealing to the induction hypothesis and the Jensen's inequality [18], we have that $E_{\pi_{s|s}}V_t^*(a_t) > V_{t-1}^*(a_t, E_{\pi_{s|s}})$, and therefore the future reward is maximized by scheduling the channel with the lowest idle age, i.e.,

$$a_t^* = \min_{a_t} \{x_t(a_t)\} \quad s.t. \quad V_{t-1}^*(\mathbf{S}_{t-1})|_{a_t^*} = \max_{a_t} \{V_{t-1}^*(\mathbf{S}_{t-1})|_{a_t}\}.$$

Finally, we proceed to verify the induction hypothesis. At $t = 2$, the optimal future reward equals the optimal immediate reward at the horizon $t = 1$, i.e.,

$$V_1^*(\mathbf{S}_1) = R_1(\mathbf{S}_1, a_1^*) := \max_{a_1 \in \mathcal{A}_1} \{R_1(\mathbf{S}_1, a_1)\}. \quad (35)$$

Since $R_1(\mathbf{S}_1, a_1)$ is linear in the strength belief, using the property of convex function [18], we have that $V_1^*(\mathbf{S}_1)$ is convex in the strength beliefs, which validates the induction hypothesis. Then, using backward induction, we establish the proposition. \blacksquare

Remarks: Proposition 5 explicitly illustrates the effect of the PU occupancy and age on the optimal decisions. With fading state beliefs equalized and the classic “exploitation vs. exploration” tradeoff neutralized (by definition of the genie-aided system), we saw that: 1) lower idle age on the PU channel favors the immediate reward by allowing more idle time on the channel; and 2) lower idle age also favors the future reward by way of better freshness of the channel fading feedback.

Thus by studying the full-observation genie-aided system, via the results in Propositions 4 and 5, we have decomposed the tradeoffs associated with the channel occupancy and the fading state beliefs in the original system. Indeed, the results in Propositions 4 and 5 rigorously support the understanding we developed earlier in Section III on the tradeoffs in the original system.

V. NUMERICAL RESULTS & FURTHER DISCUSSIONS

In this section, we evaluate and compare the optimal rewards of the original system (denoted as V_{ori}^*) and the genie-aided system (V_{genie}^*). Also, the optimal policy in the original system is compared to a *randomized* scheduling policy (with reward denoted as V_{random}^*), where the spectrum server chooses a channel, among all the idle ones, randomly and uniformly, and allocates it to the SU for data transmission. The numerical results are collected for a two channel system with $K = 2$, horizon length $m = 6$, and discounted factor $\beta = 0.9$. For notational convenience, denote $\Delta_{ga-ori} = V_{genie}^* - V_{ori}^*$ and $\Delta_{ori-rnd} = V_{ori}^* - V_{random}^*$.

Table I records the rewards obtained under various baseline cases, for various values of $\delta \triangleq p - r$, which broadly captures the temporal memory in the channel fading. In particular, as δ decreases, the channel fading memory fades and the difference between the baselines, which are primarily differentiated by the degree to which they exploit the memory in the system, tends to decrease. This is observed from Table I.

TABLE I
COMPARISON OF REWARDS WHEN THE CHANNEL FADING MEMORY VARIES. SYSTEM PARAMETERS USED:
 $u = 1, C_I = 1, C_B = 2, x_t(1) = 10, x_t(2) = 5, \pi_{t,1}^s(1) = 0.4, \pi_{t,1}^s(2) = 0.7$

δ_{pr}	Δ_{ga-ori}	$\frac{\Delta_{ga-ori}}{V_{genie}^*} \%$	$\Delta_{ori-rnd}$	$\frac{\Delta_{ori-rnd}}{V_{ori}^*} \%$
0.8	0.0077	0.3%	0.372	14.5%
0.4	0.0068	0.29%	0.2827	9.42%
0.2	0.0023	$9.8 \times 10^{-4} \%$	0.1915	8.23%
0.1	0.0006	$2.4 \times 10^{-4} \%$	0.1836	7.93%

Next, in Table II, we compare the baseline rewards under different values of the power exponent u , used in the definitions of P_I and P_B in (1). To build a better understanding of the trend reflected in these numerical results, we consider an arbitrary mini-slot, denoted as k' , as the current mini-slot, and the following two exhaustive sets of PU occupancy histories: 1) the set of histories $h_x^I, x = 1, 2, \dots$, corresponds to the case when there are exactly x mini-slots between the current mini-slot k' and the most recent mini-slot (preceding k') when the channel was in busy state, i.e., the idle age of mini-slot k' is x ; and 2) the set of histories $h_x^B, x = 1, 2, \dots$, corresponds to the case when there are exactly x mini-slots between the current mini-slot

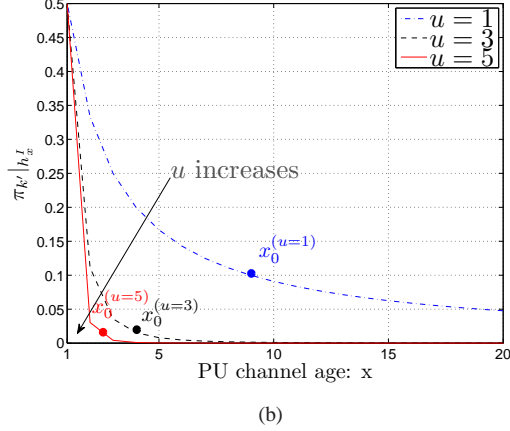
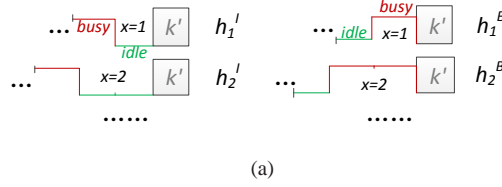


Fig. 3. Occupancy histories & the conditional idle probability.

k' and the most recent mini-slot (preceding k') when the channel was in idle state, i.e., the busy age of mini-slot k' is x . In Fig. 3(a), we plot the two sets of occupancy histories. As has been pointed out in Section II-B, the idle/busy age is a sufficient statistic for capturing the memory in the PU channels' occupancy states. Thereby, with the construction of h_x^I and h_x^B , we can examine the effect of the temporal correlation of PU occupancy on the system performance. Specifically, the conditional idle probability in mini-slot k' , given occupancy history, can be obtained as:

$$\begin{aligned}\pi_{k'}^o | h_x^I &= P_I(x) = \frac{1}{x^u + C_I}, \\ \pi_{k'}^o | h_x^B &= 1 - P_B(x) = 1 - \frac{1}{x^u + C_B}.\end{aligned}\quad (36)$$

As an example, we plot $\pi_{k'}^o | h_x^I$ in Fig. 3(b). It is clear that as u increases, the conditional probability curves become steeper. Define the threshold point, x_0 , such that for all $x > x_0$, the difference in $\pi_{k'}^o | h_x^I$ is insignificant (below 10^{-2}). Now, note from the figure that the threshold x_0 decreases with increasing u , i.e., $x_0^{(u=5)} < x_0^{(u=3)} < x_0^{(u=1)}$. This indicates that the impact of different occupancy histories on the current PU occupancy state diminishes with increasing u and thus a decreased memory in the PU occupancy. Similar argument holds when considering h_x^B . In short, the exponent u broadly captures the PU occupancy memory, and as its value increases, the memory decreases and thus the rewards under various cases, as expected, come closer with increasing u . This is illustrated in Table II.

TABLE II
COMPARISON OF REWARDS WHEN THE OCCUPANCY MEMORY VARIES. SYSTEM PARAMETERS USED:
 $p = 0.9, r = 0.1, C_I = 1, C_B = 2, x_t(1) = 0, x_t(2) = 1, \pi_{t,1}^s(1) = 0.4, \pi_{t,1}^s(2) = 0.7$

u	Δ_{ga-ori}	$\frac{\Delta_{ga-ori}}{V_{genie}^*} \%$	$\Delta_{ori-rnd}$	$\frac{\Delta_{ori-rnd}}{V_{ori}^*} \%$
1	0.0088	0.35%	0.3273	12.66%
3	0.006	0.26%	0.2201	9.59%
5	0.0043	0.2%	0.1839	8.28%

Finally, as can be seen from both tables, the original system performs very close to the genie-aided system, while the cost of measuring and sending the channel fading feedback is only $\frac{1}{N}$ of the latter. Also, the optimal policy significantly outperforms the randomized policy, indicating that the temporal correlation structure in the channel fading and PU occupancy can greatly benefit the opportunistic spectrum scheduling, when appropriately exploited.

VI. CONCLUSIONS

In this work, we studied opportunistic spectrum access for a single SU in a CR network with multiple PU channels. We formulated the problem as a partially observable Markov decision process, and examined the intricate tradeoffs in the optimal scheduling process, when incorporating the temporal correlation in *both* the channel fading and PU occupancy states. We modeled the channel fading variation with a two-state first-order Markov chain. The temporal correlation of PU traffic was modeled using age and a class of monotonically decreasing functions, which can have a long memory. The optimality of the simple greedy policy was established under certain conditions. For the general case, we individually studied the tradeoffs in the immediate reward, and the total reward. Further, by developing a genie-aided system with full observation of the channel fading feedbacks, we decomposed and characterized the multiple tiers of the intricate tradeoffs in the original system. Finally, we numerically studied the performance of the two systems and showed that the original system achieved an optimal total reward very close (within 1%) to that of the genie-aided system. Further, the optimal policy in the original system significantly outperformed randomized scheduling, highlighting the merit of exploiting the temporal correlation in the system states. We believe that our formulation and the insights we have obtained open up new horizons in better understanding spectrum allocation in cognitive radio networks, with problem settings that go beyond the traditional ones.

REFERENCES

- [1] Q. Zhao and B. M. Sadler, "A survey of dynamic spectrum access: signal processing, networking, and regulatory policy," *IEEE Signal Processing Magazine*, pp. 79–89, May 2008.
- [2] I. F. Akyildiz, W.-Y. Leea, M. C. Vuran, and S. Mohantya, "Next generation/dynamic spectrum access/cognitive radio wireless networks: a survey," *Computer Networks*, vol. 50, no. 13, pp. 2127–2159, September 2006.
- [3] B. Fette, *Cognitive Radio Technology*. Academic Press, 2006.
- [4] J. Mitola III, "Cognitive radio: an integrated agent architecture for software defined radio," Ph.D. dissertation, KTH Royal Institute of Technology, 2000.
- [5] Q. Zhang and S. Kassam, "Finite-state markov model for rayleigh fading channels," *IEEE Transactions on Communications*, vol. 47, no. 11, pp. 1688 – 1692, November 1999.
- [6] E. N. Gilbert, "Capacity of burst-noise channels," *The Bell System Technical Journal*, pp. 1253–1265, September 1960.
- [7] Q. Zhao, B. Krishnamachari, and K. Liu, "On myopic sensing for multi-channel opportunistic access: structure, optimality, and performance," *IEEE Transactions on Wireless Communications*, vol. 7, no. 12, pp. 5431–5440, 2004.
- [8] S. Murugesan, P. Schniter, and N. B. Shroff, "Multiuser scheduling in a markov-modeled downlink using randomly delayed arq feedback," Tech. Rep., 2011.
- [9] K. Kar, X. Luo, and S. Sarkar, "Throughput-optimal scheduling in multichannel access point networks under infrequent channel measurements," in *Proc. INFOCOM'07*, May 2007.
- [10] L. Johnston and V. Krishnamurthy, "Opportunistic file transfer over a fading channel: A pomdp search theory formulation with optimal threshold policies," *IEEE Transactions on Wireless Communications*, vol. 5, no. 2, pp. 394–405, 2006.
- [11] X. Xiao, K. Liu, and Q. Zhao, "Opportunistic spectrum access in self similar primary traffic," *EURASIP Journal on Advances in Signal Processing: Special Issue on Dynamic Spectrum Access for Wireless Networking (invited)*, March 2009.
- [12] Y. Chen, Q. Zhao, and A. Swami, "Distributed spectrum sensing and access in cognitive radio networks with energy constraint," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 783–797, February 2009.
- [13] F. C. Commission, "FCC frees up vacant TV airwaves for 'super wi-fi' technologies," FCC Spectrum Policy Task Force, Tech. Rep., September 2010.
- [14] I. Akyildiz, W.-Y. Lee, M. Vuran, and S. Mohanty, "A survey on spectrum management in cognitive radio networks," *IEEE Communications Magazine*, vol. 46, no. 4, pp. 40–48, April 2008.
- [15] R. D. Smallwood and E. J. Sondik, "The optimal control of partially observable markov processes over a finite horizon," *Operations Research*, vol. 21, no. 5, pp. 1071–1088, September 1973.
- [16] S. Geirhofer, L. Tong, and B. M. Sadler, "A measurement-based model for dynamic spectrum access in WLAN channels," in *Proc. IEEE Military Communications Conference (MILCOM)*, October 2006.
- [17] D. Willkomm, S. Machiraju, J. Bolot, and A. Wolisz, "Primary users in cellular networks: A large-scale measurement study," in *Proc. IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, October 2008, pp. 1–11.
- [18] B. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.